

REMARKS

Claims 1-8 were presented for examination and were pending in this application. In the latest Office Action, claims 1-8 were rejected. On the basis of the following remarks, consideration of this application and allowance of all pending claims are requested.

Claims 1 and 3-8 were rejected as anticipated by U.S. Patent No. 6,741,572 to Graves et al. Claim 2 was rejected as made obvious by Graves in view of U.S. Patent No. 6,246,692 to Dai et al. Applicant respectfully traverses the rejections.

The claims recite a network system for interconnecting a set of packet-switching network elements. The claimed network system comprises a set of nodes that interface with the packet-switching network elements, where the nodes are connected to each other by variable capacity connections. In each connection, data are transported from a source node to a destination node, and the connection has a capacity and a traffic load. In the claimed network system, the capacity of each connection is controlled from its destination node, and that control is based at least in part on the traffic loads associated with the connections configured to transport data to the same destination node. Graves does not disclose (or suggest, in combination with Dai) the claimed network system for a number of reasons, at least some of which are provided below.

“connection of variable capacity”

As recited in claim 1, each node provides a connection of variable capacity to the other nodes of the network system. This variable capacity allows for adjustment of the use of the network based on need, thereby improving the utilization of capacity for user data traffic. The network described in Graves does not have this claimed feature.

In the Office Action, the examiner argues that the lambda switches in Graves are the claimed nodes; however, these switches receive incoming traffic from each via other optical wavelengths channels, which have constant capacities (of 2.5 Gbps for OC-48 or 10 Gbps for OC-192, for example) instead of variable capacities. (col. 10, lines 60-67.)

As Graves explains, the “configuration controllers send transmission capacity request messages to configuration controllers at other nodes on the possible path and receive reply messages The reply messages may indicate whether the transmission channel can be allocated.” (col. 3, lines 15-21.) Graves further explains that “transmission channel reallocation is performed by negotiation between nodes.” (col. 5, lines 15-16.) Instead of providing connections of variable capacities, therefore, Graves is limited to the reallocation of channels of constant capacity. Importantly, Graves does not disclose varying the capacities of these channels, so the network system in Graves does not have connections of variable capacities.

The difference between Graves technique of reallocating channels of constant capacity is fundamentally different than providing a connection of variable capacity, and it leads to a number of deficiencies in Graves. For example, the reallocation of constant bandwidth channels in Graves requires that the associated network capacity be taken out of service. (col. 3, lines 27-28; col. 10, lines 6-7, 48-49; col. 11, lines 25-28, 49-50; col. 14, lines 18-19; col. 16, lines 31, 41-42.) Channel reallocation requires specifically “requesting,” “discovering,” “setting up,” “reserving,” “booking,” “implementing,” and “bringing to service” any channels being reallocated before those channels can transport user traffic. Because the channels are not carrying traffic during this reallocation or non-allocation, the associated network capacity is not serving the purpose of the communications network.

In contrast to Graves, the claimed invention dynamically optimizes the capacities of variable capacity connections. Unlike with reallocation, this technique can be performed while the variable capacity connections are transporting user data. The claimed invention, therefore, provides improved utilization of network capacity for user data traffic over the network system of Graves.

“the capacity of each connection is controlled from its destination node”

Claim 1 recites that the capacity of each connection is controlled from its destination node. In this way, a single node — the destination node — controls the variable capacity of the connection for any given connection between two nodes.

As explained above, Graves does not provide a system that has connections of variable capacity, so Graves cannot disclose that this capacity is controlled by a destination node. Moreover, the channel allocation process in Graves is also not controlled by a destination node. Instead, the control of channel allocation in Graves involves multiple nodes, where the “transmission channel reallocation is performed by negotiation between nodes.” (col. 5 lines 15-16.) This is contrasted with the claimed system, where the capacity of each connection is controlled from its single destination node.

A network system such as the claimed system in which the connection capacities are controlled by their destination nodes has a less complex and thus less costly implementation. Such system also has a deterministic performance that does not deteriorate regardless of the traffic loads on the network. Because Graves requires several nodes to be involved in the process, the system is more complicated. Moreover, in Graves a time and processing power consuming multi-stage inter-node messaging process is required to control channels in the

network. In fact, each node in Graves requires a “contract manager” and “configuration controller” with one or more microprocessors and associated buses, interfaces, and storage devices to handle the messaging and processing load. (col. 4, lines 9-10; FIG 5B.) Since the messaging and processing per Graves is software based, it has inherently non-deterministic latency, which generally worsens when the demand for message processing increases. The negotiation process among multiple nodes required to reallocate channels, while increasing the cost and complexity of the network via the required processing capacity, does not provide deterministic channel allocation completion times and hence the performance of such network is non-deterministic.

In contrast, the claimed invention controls connection capacities from the corresponding destination nodes. This allows optimization of the capacities of the connections transporting data to it cyclically based on data loads of the connections, with a fixed optimization process cycle time. This provides fixed connection capacity control process completion time regardless of the volume of connection capacity changes demanded by the data load patterns of the connections. Therefore, having the single destination node control the capacity of each connection in the network, provides an optimized and deterministic network performance under all traffic load variation scenarios.

“... based at least in part on the traffic loads associated with the connections”

In addition to having the destination node control the variable capacity of each connection, claim 1 recites that this control is “based at least in part on the traffic loads associated with the connections configured to transport data to that destination node.” This, too, is not disclosed in Graves.

Graves discloses that “if the software run by the TAP 2511 determines that the buffer fill for a particular port is large, and the rate of fill for that buffer is also large and positive, the software may generate a request to allocate additional transmission capacity.” (col. 12, lines 63-67.) Graves further explains that “[t]he software running on the LNP 2512 formulates and sends wavelength channel request messages to CMs 25 along the preferred paths,” (col. 14, lines 18-20), and that “[t]he CM 25C at node C then sends a wavelength channel request message s2 to the CM 25D at node D, specifying the source node (i.e. node C) the other intermediate nodes in the path (i.e. node E) and the destination node (i.e. Node A),” (col. 15, lines 9-13). Moreover, in Graves, the channel request based channel booking and reallocation is a multi-stage processing and messaging process involving software programs running on multiple nodes. (col. 12, line 41, to col. 17, line 19.) Therefore, this reallocation process does not have a deterministic completion time. Graves further explains, “A booked wavelength channel stays booked unless and until a subsequent reconfiguration activity allocates the wavelength channel, i.e. the booked status is a quasi-permanent status.” (col. 14, lines 50-53.) This statement is an admission that ‘booked’ channels within the network in Graves are quasi-permanent — i.e., they stay allocated between their given source and destination router, regardless of traffic loads across the network, until reallocated based on another request and only after the resulting multi-node, multi-stage messaging and processing with a non-deterministic completion time.

It can be appreciated from the foregoing that the wavelength channel requests, based on which the channels are reallocated in Graves, do not reflect traffic loads. Instead, they suggest only that at some non-deterministic amount of time earlier there was not sufficiently capacity between certain routers, and that, consequently, a channel request message was formulated and sent. At the time an additional wavelength channel can get allocated based on the request in

Graves, the actual traffic load that would demand an additional channel may not exist, since packet traffic capacity demands are not predictable but instead can change completely unpredictably. Accordingly, by the time an additional channel is allocated based on a channel request between a pair of routers in Graves, the actual traffic load between them may have decreased. This traffic load may decrease to a level at which, to achieve higher network data throughput, network capacity allocated between that pair of routers should instead be reduced to allow increasing the network capacity allocated between a different pair of routers, which, at that time, do demand more network capacity based on the actual existing traffic load level between them.

The claimed invention controls the capacities based on their traffic loads, not based on channel requests as in Graves. Unlike in Graves, therefore, the claimed invention can maximize the data throughput of the network when connection capacities are optimized based on their actual data loads, instead of based on requests for additional capacity generated a non-deterministic time before the additional capacity could be allocated. Accordingly, Graves does not disclose or suggest this claimed feature.

For at least the foregoing reasons, Graves does not anticipate claim 1, or any of claims 3-8, which depend from claim 1. Because the obviousness rejection of claim 2 applied Graves in the same way that the anticipation rejection of claim 1 applied Graves, claim 2 is patentable over the combination of Graves and Dai for the same reasons set forth above.

Based on the foregoing, the application is in condition for allowance of all claims, and a Notice of Allowance is respectfully requested. If the examiner believes for any reason direct

contact would help advance the prosecution of this case to allowance, the examiner is encouraged to telephone the undersigned at the number given below.

Respectfully submitted,
MARK SANDSTROM

Dated: April 12, 2006

By: /Robert A. Hulse/

Robert A. Hulse, Reg. No. 48,473
Attorney for Applicant
Fenwick & West LLP
801 California Street
Mountain View, CA 94041
Tel.: (415) 875-2444
Fax: (415) 281-1350